Distributed Large Scale Network Utility Maximization

Danny Bickson, Yoav Tock
IBM Haifa Research Lab
Mount Carmel, Haifa 31905, Israel
Email: {dannybi,tock}@il.ibm.com

Argyris Zymnis, Stephen P. Boyd
Department of Electrical Engineering
Stanford University
350 Serra Mall, Packard 243,
Stanford CA 94305, USA
Email: {azymnis,boyd}@stanford.edu

Danny Dolev
School of Computer Science and Engineering
Hebrew University of Jerusalem
Jerusalem 91904, Israel
Email: dolev@cs.huji.ac.il

arXiv:0901.2684v2 [cs.IT] 9 May 2009

***Abstract*— Recent work by Zymnis *et al.* proposes an efficient primal-dual interior-point method, using a truncated Newton method, for solving the network utility maximization (NUM) problem. This method has shown superior performance relative to the traditional dual-decomposition approach. Other recent work by Bickson *et al.* shows how to compute efficiently and distributively the Newton step, which is the main computational bottleneck of the Newton method, utilizing the Gaussian belief propagation algorithm.**

In the current work, we combine both approaches to create an efficient distributed algorithm for solving the NUM problem. Unlike the work of Zymnis, which uses a centralized approach, our new algorithm is easily distributed. Using an empirical evaluation we show that our new method outperforms previous approaches, including the truncated Newton method and dual-decomposition methods. As an additional contribution, this is the first work that evaluates the performance of the Gaussian belief propagation algorithm vs. the preconditioned conjugate gradient method, for a large scale problem.

I. Introduction

We consider a network that supports a set of flows, each of which has a nonnegative flow rate, and an associated utility function. Each flow passes over a route, which is a subset of the edges of the network. Each edge has a given capacity, which is the maximum total traffic (the sum of the flow rates through it) it can support. The network utility maximization (NUM) problem is to choose the flow rates to maximize the total utility, while respecting the edge capacity constraints [1], [2]. We consider the case where all utility functions are concave, in which case the NUM problem is a convex optimization problem.

A standard technique for solving NUM problems is based on dual decomposition [3], [4]. This approach yields fully decentralized algorithms, that can scale to very large networks. Dual decomposition was first applied to the NUM problem in [5], and has led to an extensive body of research on distributed algorithms for network optimization [6]–[8] and new ways to interpret existing network protocols [9].

Recent work by Zymnis *et al.* presented a specialized primal-dual interior-point method for the NUM problem [10]. Each Newton step is computed using the preconditioned conjugate gradient method (PCG). This proposed method had a significant performance improvement over the dual decomposition approach, especially when the network is congested. Furthermore, the method can handle utility functions which are not strictly concave. The main drawback of the primal-dual method is that it is centralized, while the dual decomposition methods are easily distributed.

Other recent work by Bickson *et al.* [11] proposes an efficient way for computing the Newton step, which is the main computational effort of the primal-dual interior-point method using the Gaussian belief propagation (GaBP) algorithm, which is an efficient distributed algorithm.

In the current paper we propose to combine both previous approaches. We present an efficient primal-dual interior point method, where the Newton step computed in each iteration is computed using the GaBP algorithm. Using extensive simulations with very large scale networks we compare the performance of our novel method to previous approaches including an interior-point method using PCG, and dual decomposition methods. Despite of being distributed, our new construction exhibits significant performance improvements over previous approaches.

Furthermore, we provide the first comparison of performance of the GaBP algorithm vs. the PCG method. The PCG method is a state-of-the-art method used extensively in large-scale optimization applications. Examples include ℓ_1-regularized logistic regression [12], gate sizing [13], and slack allocation [14]. Empirically, the GaBP algorithm is immune to numerical problems with typically occur in the PCG method, while demonstrating a faster convergence. The only previous work comparing the performance of GaBP vs. PCG we are aware of is [15], which used a small example of 25 nodes, and the work of [16] which used a grid of 25×25 nodes.

We believe that our approach is general and not limited to the NUM problem. It could potentially be used for the solution of other large scale distributed optimization problems.

This paper is organized as follows. Section II briefly overviews the NUM problem formulation. Section III outlines previous algorithms for solving the NUM problem, including dual descent and truncated Newton method. We present our new construction which utilizes the GaBP algorithm in Section IV. Section V provides simulation results comparing the performance of the GaBP based algorithm with the previous approaches. We conclude in Section VI.

II. Problem formulation

There are n flows in a network, each of which is associated with a fixed route, *i.e.*, some subset of m links. Each flow has a nonnegative *rate*, which we denote $f_1, \ldots, f_n$. With the flow j we associate a utility function $U_j : \mathbf{R} \to \mathbf{R}$, which is concave and twice differentiable, with $\mathbf{dom}\, U_j \subseteq \mathbf{R}_+$. The utility derived by a flow rate f_j is given by $U_j(f_j)$. The total utility associated with all the flows is then $U(f) = U_1(f_1) + \cdots + U_n(f_n)$.

The total traffic on a link in the network is the sum of the rates of all flows that utilize that link. We can express the link traffic compactly using the *routing* or *link-route* matrix $R \in \mathbf{R}^{m\times n}$, defined as

$$R_{ij} = \begin{cases} 1 & \text{flow } j\text{'s route passes over link } i \\ 0 & \text{otherwise.} \end{cases}$$

Each link in the network has a (positive) *capacity* $c_1, \ldots, c_m$. The traffic on a link cannot exceed its capacity, *i.e.*, we have $Rf \leq c$, where $\leq$ is used for componentwise inequality.

The NUM problem is to choose the rates to maximize total utility, subject to the link capacity and the nonnegativity constraints:

$$\begin{array}{ll} \text{maximize} & U(f) \\ \text{subject to} & Rf \leq c, \quad f \geq 0, \end{array} \tag{1}$$

with variable $f \in \mathbf{R}^n$. This is a convex optimization problem and can be solved by a variety of methods. We say that f is *primal feasible* if it satisfies $Rf \leq c$, $f \geq 0$.

The dual of problem (1) is

$$\begin{array}{ll} \text{minimize} & \lambda^T c + \sum_{j=1}^n (-U_j)^*(-r_j^T\lambda) \\ \text{subject to} & \lambda \geq 0, \end{array} \tag{2}$$

where $\lambda \in \mathbf{R}^m_+$ is the dual variable associated with the capacity constraint of problem (1), r_j is the jth column of R and $(-U_j)^*$ is the conjugate of the negative jth utility function [17, §3.3],

$$(-U_j)^*(a) = \sup_{x\geq 0}(ax + U_j(x)).$$

We say that λ is *dual feasible* if it satisfies $\lambda \geq 0$ and $\lambda \in \cap_{j=1}^n \mathbf{dom}(-U_j)^*$.

III. Previous work

In this section we give a brief overview of the dual-decomposition method and the primal-dual interior point method proposed in [10].

A. Dual decomposition

Dual decomposition [3]–[6] is a projected (sub)gradient algorithm for solving problem (2), in the case when all utility functions are strictly concave. We start with any positive λ, and repeatedly carry out the update

$$\begin{aligned} f_j &:= \operatorname*{argmax}_{x\geq 0} \left(U_j(x) - x(r_j^T\lambda)\right), \quad j = 1, \ldots, n, \\ \lambda &:= \left(\lambda - \alpha\left(c - Rf\right)\right)_+, \end{aligned}$$

where $\alpha > 0$ is the step size, and x_+ denotes the entrywise nonnegative part of the vector x. It can be shown that for small enough α, f and λ will converge to $f^\star$ and $\lambda^\star$, respectively, provided all U_j are differentiable and strictly concave. The term $s = c - Rf$ appearing in the update is the *slack* in the link capacity constraints (and can have negative entries during the algorithm execution). It can be shown that the slack is exactly the gradient of the dual objective function.

Dual decomposition is a distributed algorithm. Each flow is updated based on information obtained from the links it passes over, and each link dual variable is updated based only on the flows that pass over it.

B. Primal-dual interior point method

The primal-dual interior-point method is based on using a Newton step, applied to a suitably modified form of the optimality conditions. The modification is parametrized by a parameter t, which is adjusted during the algorithm based on progress, as measured by the actual duality gap (if it is available) or a surrogate duality gap (when the actual duality gap is not available).

We first describe the search direction. We modify the complementary slackness conditions to obtain the modified optimality conditions

$$\begin{aligned} -\nabla U(f) + R^T\lambda - \mu &= 0 \\ \mathbf{diag}(\lambda)s &= (1/t)\mathbf{1} \\ \mathbf{diag}(\mu)f &= (1/t)\mathbf{1}, \end{aligned}$$

where $t > 0$ is a parameter that sets the accuracy of the approximation. (As $t \to \infty$, we recover the optimality conditions for the NUM problem.) Here we implicitly assume that $f, s, \lambda, \mu > 0$. The modified optimality conditions can be compactly written as $r_t(f, \lambda, \mu) = 0$, where

$$r_t(f, \lambda, \mu) = \begin{bmatrix} -\nabla U(f) + R^T\lambda - \mu \\ \mathbf{diag}(\lambda)s - (1/t)\mathbf{1} \\ \mathbf{diag}(\mu)f - (1/t)\mathbf{1} \end{bmatrix}.$$

The primal-dual search direction is the Newton step for solving the nonlinear equations $r_t(f, \lambda, \mu) = 0$. If $y = (f, \lambda, \mu)$ denotes the current point, the Newton step $\Delta y = (\Delta f, \Delta\lambda, \Delta\mu)$ is characterized by the linear equations

$$r_t(y + \Delta y) \approx r_t(y) + r_t'(y)\Delta y = 0,$$

which, written out in more detail, are

$$\begin{bmatrix} -\nabla^2 U(f) & R^T & -I \\ -\mathbf{diag}(\lambda)R & \mathbf{diag}(s) & 0 \\ \mathbf{diag}(\mu) & 0 & \mathbf{diag}(f) \end{bmatrix} \begin{bmatrix} \Delta f \\ \Delta\lambda \\ \Delta\mu \end{bmatrix} = -r_t(f, \lambda, \mu). \tag{3}$$

During the algorithm, the parameter t is increased, as the primal and dual variables approach optimality. When we have easy access to a dual feasible point during the algorithm, we can make use of the exact duality gap η to set the value of t; in other cases, we can use the surrogate duality gap $\hat{\eta}$.

The primal-dual interior point algorithm is given in [17, §11.7], [18].

The most expensive part of computing the primal-dual search direction is solving equation (3). For problems of modest size, *i.e.*, with m and n no more than 10^4, it can be solved using direct methods such as a sparse Cholesky decomposition.

For larger problem instances [10] proposes to solve (3) *approximately*, using a preconditioned conjugate gradient (PCG) algorithm [19, §6.6], [20, chap. 2], [21, chap. 5]. When an iterative method is used to approximately solve a Newton system, the algorithm is referred to as an *inexact, iterative*, or *approximate* Newton method (see [20, chap. 6] and its references). When an iterative method is used inside a primal-dual interior-point method, the overall algorithm is called a *truncated-Newton primal-dual interior-point method*. For details of the PCG algorithm, we refer the reader to the references cited above. Each iteration requires multiplication of the matrix by a vector, and a few vector inner products.

IV. Our new construction

Previous work of Zymnis *et al.* [10] shows that when applying the interior-point Newton method to the NUM problem, each Newton step involves a solution of Eq. 3, where the solution $(\Delta f, \Delta\lambda, \Delta\mu)^T$ is the Newton search direction.

Recent results by Bickson *et al.* [22], [23] utilizes the GaBP algorithm as an efficient distributed algorithm for solving a system of linear equations. For utilizing the GaBP algorithm, we first normalize Eq. (3) by $(1, -1/\lambda, -1/\mu)$ to get the following equivalent system of linear equations:

$$\begin{bmatrix} -\nabla^2 U(f) & R^T & -I \\ R & -\mathbf{diag}(s/\lambda) & 0 \\ -I & 0 & \mathbf{diag}(f/\mu) \end{bmatrix} \begin{bmatrix} \Delta f \\ \Delta\lambda \\ \Delta\mu \end{bmatrix} = -\hat{rt}(f,\lambda,\mu). \tag{4}$$

where $\hat{rt}(f,\lambda,\mu) = ((-\nabla U(f) + R^T - \mu)^T, (-s + (\lambda/t))^T, (-f + (\mu/t))^T)^T$. Note that the new system of linear equations is symmetric, a condition required by the GaBP algorithm.

The formulation (4) allows us to shift the linear system of equations from an algebraic to a probabilistic domain. Instead of solving a deterministic vector-matrix linear equation, we now solve an inference problem in a graphical model describing a certain Gaussian distribution function. Following [24] we define the joint covariance matrix

$$A \triangleq \begin{bmatrix} -\nabla^2 U(f) & R^T & -I \\ R & -\mathbf{diag}(s/\lambda) & 0 \\ -I & 0 & \mathbf{diag}(f/\mu) \end{bmatrix} \tag{5}$$

and the shift vector $b \triangleq ((-\nabla U(f) + R^T - \mu)^T, (-s + (\lambda/t))^T, (-f + (\mu/t))^T)^T$. We further denote the search direction $x = (\Delta f^T, \Delta\lambda^T, \Delta\mu^T)^T$.

Given the covariance matrix A and the shift vector b, one can write explicitly the Gaussian density function

$$p(x) \sim \exp(-1/2 x^T A x + b^T x)$$

Now, we are interested in computing the MAP assignment:

$$x^* = \underset{x}{\operatorname{argmax}} \exp(-1/2 x^T A x + b^T x)$$

The corresponding graph of the covariance matrix A is $\mathcal{G}$, with edge potentials ('compatibility functions') ψ_{ij} and self-potentials ('evidence') ϕ_i. These graph potentials are determined according to the following pairwise factorization of the Gaussian distribution $p(x) \propto \prod_{i=1}^n \phi_i(x_i) \prod_{\{i,j\}} \psi_{ij}(x_i, x_j)$, resulting in $\psi_{ij}(x_i, x_j) \triangleq \exp(-x_i A_{ij} x_j)$, and $\phi_i(x_i) \triangleq \exp\left(b_i x_i - A_{ii} x_i^2/2\right)$. The set of edges $\{i, j\}$ corresponds to the set of non-zero entries in A (Eq. 5). Hence, we would like to calculate the marginal densities, which must also be Gaussian,

$$p(x_i) \propto \mathcal{N}(\mu_i = \{A^{-1}b\}_i, P_i^{-1} = \{A^{-1}\}_{ii}),$$

where μ_i and P_i are the marginal mean and inverse variance (a.k.a. precision), respectively. Recall that, according to [24], the inferred mean μ is identical to the desired solution of (Eq. 4). The GaBP update rules are summarized in Table I. We use the notation $\mathbb{N}(i)$ as the set of node i graph neighbors, excluding i.

#	Stage	Operation
1.	*Initialize*	Compute $P_{ii} = A_{ii}$ and $\mu_{ii} = b_i/A_{ii}$. Set $P_{ki} = 0$ and $\mu_{ki} = 0$, $\forall k \in \mathbb{N}(i)$.
2.	*Iterate*	Propagate P_{ki} and μ_{ki}, $\forall k \in \mathbb{N}(i)$. Compute $P_{i\setminus j} = P_{ii} + \sum_{k\in\mathbb{N}(i)\setminus j} P_{ki}$ $\mu_{i\setminus j} = P_{i\setminus j}^{-1}(P_{ii}\mu_{ii} + \sum_{k\in\mathbb{N}(i)\setminus j} P_{ki}\mu_{ki})$. Compute $P_{ij} = -A_{ij} P_{i\setminus j}^{-1} A_{ji}$ $\mu_{ij} = -P_{ij}^{-1} A_{ij} \mu_{i\setminus j}$.
3.	*Check*	If P_{ij} and μ_{ij} did not converge, return to #2. Else, continue to #4.
4.	*Infer*	$P_i = P_{ii} + \sum_{k\in\mathbb{N}(i)} P_{ki}$ $\mu_i = P_i^{-1}(P_{ii}\mu_{ii} + \sum_{k\in\mathbb{N}(i)} P_{ki}\mu_{ki})$.
5.	*Output*	$x_i = \mu_i$

TABLE I: Computing $x = A^{-1}b$ via GaBP [23].

It is known that if GaBP converges, it results in exact inference [25]. Determining the exact region of convergence remain open research problems. All that is known is a sufficient (but not necessary) condition stating that GaBP converges when the spectral radius satisfies $\rho(|I_K - A|) < 1$ [26], [27]. A stricter sufficient condition [25], determines that the matrix A must be diagonally dominant (*i.e.*, $|A_{ii}| > \sum_{j\neq i} |A_{ij}|, \forall i$) in order for GaBP to converge. Recently, a new technique for forcing convergence for any column-dependent matrices is proposed in [28]. An upper bound on convergence speed is given in [11].

V. Experimental results

A. Small experiment

In our first example we look at the performance of our method on a small network. The utility functions are all logarithmic, *i.e.*, $U_j(f_j) = \log f_j$. There are $n = 10^3$ flows, and $m = 2 \cdot 10^3$ links. The elements of R are chosen randomly and independently, so that the average route length is 10 links. The link capacities c_i are chosen independently from a uniform

distribution on $[0.1, 1]$. For this particular example, there are about 10^4 nonzero elements in R (0.5% density).

We compare three different algorithms for solving the NUM problem: The dual-decomposition method, a truncated Newton method via PCG and a customized Newton method via the GaBP solver. Out of the examined algorithms, the Newton method is centralized, while the dual-decomposition and GaBP solver are distributed algorithms. The source code of our Matlab simulation is available on [29].

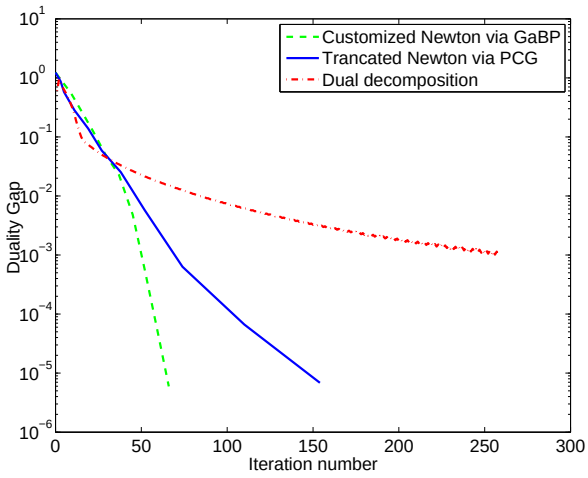

Fig. 1: Convergence rate using the small settings.

Figure 1 depicts the solution quality, where the X-axis represents the number of algorithm iterations, and the Y-axis is the surrogate duality gap (using a logarithmic scale). As clearly shown, the GaBP algorithm has a comparable performance to the sparse Cholesky decomposition, while it is a distributed algorithm. The dual decomposition method has much slower convergence.

B. Larger experiment

Our second example is too large to be solved using the primal-dual interior-point method with direct search direction computation, but is readily handled by the truncated-Newton primal-dual algorithm using PCG, the dual decomposition method and the customized Newton method via GaBP. The utility functions are all logarithmic: $U_j(f_j) = \log f_j$. There are $n = 10^4$ flows, and $m = 2 \cdot 10^4$ links. The elements of R and c are chosen as for the small example. For dual decomposition, we initialized all λ_i as 1. For the interior-point method, we initialized all λ_i and μ_i as 1. We initialize all f_j as γ, where we choose γ so that $Rf \leq 0.9c$.

Our experimental results shows, that as the system size grows larger, the GaBP solver has favorable performance. Figure 2 plots the duality gap of both algorithms, vs. the number of iterations performed.

Figure 3 shows that in terms of Newton steps, both methods had comparable performance. The Newton method via the GaBP algorithm converged in 11 steps, to an accuracy of

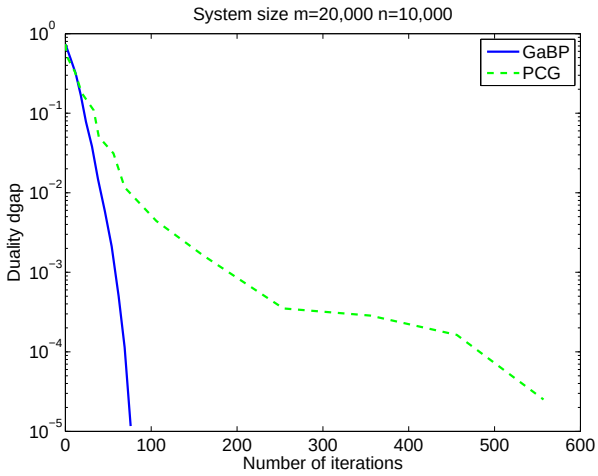

Fig. 2: Convergence rate in the larger settings.

10^{-4} where the truncated Newton method implemented via PCG converged in 13 steps to the same accuracy. However, when examining the iteration count in each Newton step (the Y-axis) we see that the GaBP remained constant, while the PCG iterations significantly increase as we are getting closer to the optimal point.

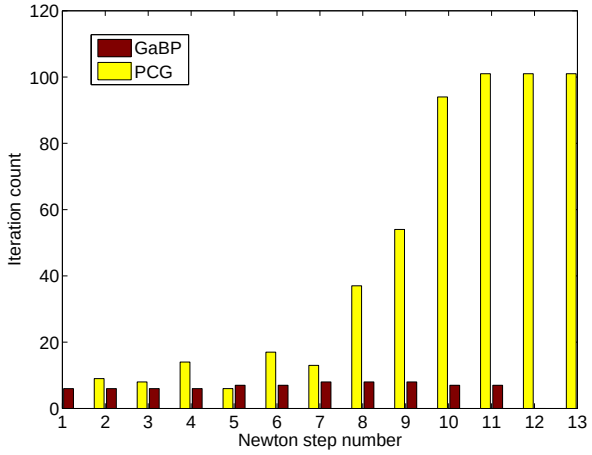

Fig. 3: Iteration count per Newton step.

We have experimented with larger settings, up to $n = 10^5$ flows, and $m = 2 \cdot 10^5$ links. The GaBP algorithm converged in 11 Newton steps with 7-9 inner iteration in each Newton step. The PCG method converged in 16 Newton steps with an average of 45 inner iterations.

C. Numerical issues

Overall, we have observed three types of numerical problems with the PCG method. First, the PCG Matlab implementation runs into numerical problems and failed to compute

the search direction. Second, the line search failed, which means that no progress is possible in the computed direction without violating the problem constraints. Third, when getting close to the optimal solution, the number of PCG iterations significantly increases.

The numerical problems of the PCG algorithm are well known, see of example [30], [31]. In contrary, the GaBP algorithm did not suffer from the above numerical problems.

Furthermore, the PCG is harder to distribute, since in each PCG iteration a vector dot product and a matrix product are performed. Those operations are global, unlike the GaBP which exploits the sparseness of the input matrix.

VI. Conclusion

We propose an efficient distributed solution of the NUM problem using a customized Newton method, implemented via the GaBP algorithm. We compare the customized Newton method performance with state-of-the-art algorithms, including a dual descent method and a truncated Newton method, over large scale settings. We observe both faster convergence of the GaBP algorithm compared to both the preconditioned conjugate gradient and sparse Cholesky factorization. Furthermore, the GaBP does not suffer from numerical problems which affect the performance of the preconditioned conjugate gradient method.

We believe that the NUM problem serves as a case study for demonstrating the superior performance of the GaBP algorithm in solving sparse systems of linear equations. Since the problem of solving a system of linear equations is a fundamental problem in computer science and engineering, we envision many other applications for our proposed method.

Acknowledgment

Danny Dolev is Incumbent of the Berthold Badler Chair in Computer Science. Danny Dolev was supported in part by the Israeli Science Foundation (ISF) Grant number 0397373.

References

[1] R. Srikant, *The Mathematics of Internet Congestion Control*. Birkhäuser, 2004.

[2] D. Bertsekas, *Network Optimization: Continuous and Discrete Models*. Athena Scientific, 1998.

[3] G. B. Dantzig and P. Wolfe, "Decomposition principle for linear programs," *Operations Research*, vol. 8, pp. 101–111, 1960.

[4] N. Z. Shor, *Minimization Methods for Non-Differentiable Functions*. Springer-Verlag, 1985.

[5] F. Kelly, A. Maulloo, and D. Tan, "Rate control for communication networks: Shadow prices, proportional fairness and stability," *Journal of the Operational Research Society*, vol. 49, pp. 237–252, 1997.

[6] S. H. Low and D. E. Lapsley, "Optimization flow control I: Basic algorithms and convergence," *IEEE/ACM Transactions on Networking*, vol. 7, no. 6, pp. 861–874, Dec. 1999.

[7] M. Chiang, S. H. Low, A. R. Calderbank, and J. C. Doyle, "Layering as optimization decomposition: A mathematical theory of network architectures," *Proceedings of the IEEE*, vol. 95, no. 1, pp. 255–312, Jan. 2007.

[8] D. Palomar and M. Chiang, "A tutorial on decomposition methods and distributed network resource allocation," *IEEE Journal of Selected Areas in Communication*, vol. 24, no. 8, pp. 1439–1451, Aug. 2006.

[9] S. H. Low, "A duality model of TCP and queue management algorithms," *IEEE/ACM Transactions on Networking*, vol. 11, no. 4, pp. 525–536, Aug. 2003.

[10] A. Zymnis, N. Trichakis, S. Boyd, and D. Oneill, "An interior-point method for large scale network utility maximization," in *Proceedings of the Allerton Conference on Communication, Control, and Computing*, 2007.

[11] D. Bickson, Y. Tock, D. Dolev, and O. Shental, "Polynomial linear programming with Gaussian belief propagation," in *the 46th Allerton Conf. on Communications, Control and Computing*, Monticello, IL, USA, 2008.

[12] K. Koh, S.-J. Kim, and S. Boyd, "An interior point method for large-scale ℓ_1-regularized logistic regression," *Journal of Machine Learning Research*, vol. 8, pp. 1519–1555, July 2007.

[13] S. Joshi and S. Boyd, "An efficient method for large-scale gate sizing," *IEEE Trans. Circuits and Systems I: Fundamental Theory and Applications*, vol. 5, no. 9, pp. 2760–2773, Nov. 2008.

[14] ——, "An efficient method for large-scale slack allocation," in *Engineering Optimization*, 2008.

[15] E. B. Sudderth, "Embedded trees: Estimation of Gaussian processes on graphs with cycles," Master's thesis, University of California at San Diego, February 2002.

[16] Y. Weiss and W. T. Freeman, "Correctness of belief propagation in Gaussian graphical models of arbitrary topology," *Neural Computation*, vol. 13, no. 10, pp. 2173–2200, 2001.

[17] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.

[18] S. J. Wright, *Primal-Dual Interior-Point Methods*. Society for Industrial and Applied Mathematics, 1997.

[19] J. Demmel, *Applied Numerical Linear Algebra*. SIAM, 1997.

[20] C. T. Kelley, *Iterative Methods for Linear and Nonlinear Equations*. SIAM, 1995.

[21] J. Nocedal and S. J. Wright, *Numerical Optimization*. Springer, 1999.

[22] D. Bickson, O. Shental, P. H. Siegel, J. K. Wolf, and D. Dolev, "Linear detection via belief propagation," in *Proc. 45th Allerton Conf. on Communications, Control and Computing*, Monticello, IL, USA, Sept. 2007.

[23] O. Shental, D. Bickson, P. H. Siegel, J. K. Wolf, and D. Dolev, "Gaussian belief propagation solver for systems of linear equations," in *IEEE Int. Symp. on Inform. Theory (ISIT)*, Toronto, Canada, July 2008.

[24] D. Bickson, O. Shental, P. H. Siegel, J. K. Wolf, and D. Dolev, "Gaussian belief propagation based multiuser detection," in *IEEE Int. Symp. on Inform. Theory (ISIT)*, Toronto, Canada, July 2008.

[25] Y. Weiss and W. T. Freeman, "Correctness of belief propagation in Gaussian graphical models of arbitrary topology," *Neural Computation*, vol. 13, no. 10, pp. 2173–2200, 2001.

[26] J. Johnson, D. Malioutov, and A. Willsky, "Walk-sum interpretation and analysis of Gaussian belief propagation," in *NIPS 05'*.

[27] D. M. Malioutov, J. K. Johnson, and A. S. Willsky, "Walk-sums and belief propagation in Gaussian graphical models," in *Journal of Machine Learning Research*, vol. 7, Oct. 2006.

[28] J. K. Johnson, D. Bickson, and D. Dolev, "Fixing convergence of Gaussian belief propagation," in *International Symposium on Information Theory (ISIT)*, Seoul, South Korea, July 2009.

[29] Gaussian Belief Propagation implementation in matlab [online] `http://www.cs.huji.ac.il/labs/danss/p2p/gabp/`.

[30] K. C. Toh, "Solving large scale semidefinite programs via an iterative solver on the augmented systems," *SIAM J. on Optimization*, vol. 14, no. 3, pp. 670–698, 2003.

[31] M. Kočvara and M. Stingl, "On the solution of large-scale SDP problems by the modified barrier method using iterative solvers," *Math. Program.*, vol. 109, no. 2, pp. 413–444, 2007.